Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

FLBEC: A Unified Framework Integrating Federated Learning, Blockchain, and Edge Computing for Privacy-Preserving Cybersecurity Governance in Next-Generation Networks

Subha Shamarukh [1] and Navila Sultana [2]
1 University of Rochester, Rochester, New York, USA, shamarukhsubha@gmail.com, https://orcid.org/0009-0000-2170-1541
2 University of Tulsa, Tulsa, OK, USA, navilasultana.research@gmail.com, https://orcid.org/0009-0004-8272-6282
Corresponding Author: Subha Shamarukh, **E-mail**: shamarukhsubha@gmail.com

| ABSTRACT
Across healthcare, telecommunications, and critical infrastructure, organizations are sitting on threat intelligence that could protect their peers and sharing almost none of it. The reason is not selfishness; it is a genuine legal and competitive constraint. Data that reveals one network's vulnerabilities also reveals its patients, its subscribers, or its operational secrets. This paper begins with that everyday reality and asks a practical question: can we build a system that enables organizations to learn from each other's threats without ever exposing the underlying data? We propose FLBEC, a framework that binds together Federated Learning (FL), Blockchain governance (BC), and Edge Computing intelligence (EC) to deliver threat detection that is simultaneously private, auditable, real-time, and ready standards. We evaluate FLBEC across five benchmark datasets, NSL-KDD, CICIDS-2017, CTU-13, UNSW-NB15, and a custom healthcare IoT dataset, measuring detection accuracy, false positive rate, AUC-ROC, inference latency, and communication overhead. We also run a systematic ablation study to isolate what each of the three components actually contributes. FLBEC reaches 96.3% detection accuracy, 3.7% false positive rate, and AUC-ROC of 0.963, clearing every single-paradigm and pairwise baseline. End-to-end response latency sits at 42 ms, and communication overhead drops to 1.8 GB per federated round. Every component passes the ablation test for independent statistical significance at $p < 0.05$. The results confirm that privacy, governance, and real-time performance are not a trade-off triangle, they become mutually reinforcing when the three paradigms are co-designed from the start. FLBEC maps directly to IEEE 802.1X, 3GPP Release 18, and NIST CSF 2.0, making it deployable in regulated industries without modification.

| KEYWORDS
Federated Learning, Blockchain Governance, Edge Computing, Cybersecurity, Privacy Preservation, 6G Networks, Homomorphic Encryption, Threat Detection, Zero-Trust Architecture, Continual Learning, IoT Security

| ARTICLE INFORMATION
ACCEPTED: 09 April 2026 **PUBLISHED:** 29 May 2026 **DOI:** 10.32996/jcsts.2026.8.7.13

1. INTRODUCTION

Imagine a hospital network security team watching an unfamiliar attack pattern propagate through their systems at 2 a.m. Three hospitals in other states encountered the same pattern eighteen months earlier. They contained it in hours. But no one told this team, because sharing the logs would have meant sharing data tied to patient records, and no compliance officer would sign off on that. This is not hypothetical. It is a description of how threat intelligence actually moves, or fails to move, through industries where data is both the most valuable asset and the most legally constrained one.

The financial scale of this fragmentation is staggering. Global cybercrime costs are projected to exceed $10.5 trillion annually by 2025 (Morgan, 2020), yet a large fraction of those losses come from attacks that were already known to someone in the ecosystem, just not to the victim. The core technical problem is a tension between two legitimate needs: organizations need to collaborate on threat intelligence to build better defenses, but they cannot move the data that would make collaboration possible.

Federated learning was invented precisely to break this deadlock. Instead of centralizing data, it centralizes only the learning process. Each organization trains a model locally and contributes gradient updates rather than raw records (McMahan et al., 2017). That is a genuine advance, but it does not solve the full problem. Gradient updates can be reverse engineered to recover sensitive records through model inversion attacks (Fredrikson et al., 2015). And even if gradients were perfectly private, federated learning alone provides no way to prove to a regulator, an auditor, or an incident response team that every participant behaved honestly. That verification is where blockchain earns its place: an immutable, cryptographically linked ledger of every governance decision, every policy enforcement outcome, every model update commitment (Haldar et al., 2026). The remaining gap is latency. Blockchain consensus and cloud-tier authorization add overhead that operational technology environments simply cannot absorb — a few hundred milliseconds is too long when an attacker is pushing commands toward an industrial actuator (Varanasi et al., 2026). Edge computing closes that gap by moving inference to the point where threats actually appear.

These three paradigms, federated learning, blockchain governance, and edge intelligence, have each been researched intensively in isolation. What has not been studied seriously is whether binding them together produces something qualitatively different from what any one achieves alone. That is the central question of this paper. Our hypothesis is that a principled co-design of all three can satisfy four requirements that individually optimized systems routinely fail to achieve simultaneously: GDPR- and CCPA-compliant privacy preservation; tamper-evident, real-time governance audit; sub-50 ms threat containment at the network edge; and alignment with IEEE 802.1X, 3GPP Release 18, and NIST CSF 2.0 standards.

That hypothesis did not emerge from theory alone. Recent empirical work across the literature provides converging motivation. Hasan et al. (2025a) show that AI-driven governance measurably reduces breach exposure in enterprise data environments, but also that the benefit depends on tight integration between the AI decision layer and the governance audit layer, a dependency FLBEC is designed to satisfy structurally. Hasan et al. (2026) argue that cybersecurity risk management must be embedded into national-scale information systems governance to deliver organizational resilience, not just incident response, a vision whose technical instantiation requires exactly the kind of blockchain-backed audit trail FLBEC provides. Orthi et al. (2025a) demonstrate that federated learning with privacy-preserving analytics is viable across distributed hospital networks but note that without a governance mechanism to certify participant behavior, healthcare regulators will not sanction cross-institutional model sharing. Siam et al. (2025a) show that transformer-based intrusion detection works on smart building IoT hardware when appropriately quantized — evidence that FLBEC's edge inference design is grounded in demonstrated deployment feasibility rather than theoretical performance.

FLBEC makes four specific technical contributions. First, a homomorphically encrypted federated aggregation protocol using CKKS-scheme encryption applied to top-k sparsified gradients, reducing per-round communication to 1.8 GB at 100 participants, a 7.4-fold reduction over unencrypted FL. Second, a smart contract governance layer that enforces GDPR, HIPAA, and sector-specific compliance policies at the point of model update, eliminating the post-hoc audit burden that Chakraborty et al. (2025) identify as a primary friction in FL-based enterprise data architectures. Third, a reinforcement learning self-healing edge agent that executes threat containment in under 42 ms for non-irreversible actions, with irreversible actions gated by blockchain-signed authorization to preserve human oversight. Fourth, a reference architecture mapped to IEEE 802.1X, 3GPP Release 18, and NIST CSF 2.0, following the 6G-ready design principles demonstrated by Mahin et al. (2026) and the digital twin communication architecture of Varanasi et al. (2026).

The paper is organized as follows. Section 2 surveys related literature and situates FLBEC within it. Section 3 develops the framework's theoretical foundations. Section 4 describes architectural design. Section 5 presents the experimental methodology. Section 6 reports and interprets results. Section 7 addresses open challenges. Section 8 concludes.

2. RELATED WORK

2.1 Federated Learning: From Privacy Promise to Governance Reality

When McMahan et al. (2017) introduced federated learning, the appeal was straightforward: collaborative model training without data centralization. In the years since, the security community has been working through the gap between that conceptual promise and practical deployment. Das et al. (2025a) take a significant step by combining homomorphic encryption with FL, demonstrating that gradient-level privacy leakage can be reduced to formally bounded levels with training overhead below 4× relative to unencrypted FL. Their work establishes the performance envelope within which FLBEC's encryption layer must operate. In the healthcare domain specifically, Orthi et al. (2025a) deploy FL with privacy-preserving analytics across distributed hospital systems, reporting that cross-institutional model generalization improves substantially when non-IID data distribution is explicitly accounted for during aggregation — a finding that directly shapes how FLBEC partitions its healthcare IoT evaluation data.

Das et al. (2025b) push FL into the operational technology world, designing a critical infrastructure federation that spans multiple sectors and covers 87% of MITRE ATT&CK for ICS techniques. Their honest reckoning with what FL alone cannot do, governance, real-time response, standards certification, is one of the direct motivations for FLBEC's co-design approach. Chakraborty et al. (2025) explore the FL-blockchain combination in a data lakehouse context, showing that the two paradigms complement each other naturally. Their central insight, that post-hoc compliance auditing creates too much operational friction to be sustainable, is something FLBEC takes seriously as a design requirement, not just an observation. The homomorphic encryption benchmarking of Mohonta et al. (2026), which establishes CKKS as the best-performing scheme for floating-point gradient encryption, gives FLBEC's protocol selection an empirically grounded basis rather than an arbitrary one.

2.2 Blockchain as an Active Governance Layer

What makes recent blockchain-for-security work interesting is a shift in how blockchain is being used. Earlier work treated blockchain primarily as a record-keeping system, a better audit log. More recent work treats it as an active governance layer: policy execution, access control enforcement, compliance obligation fulfillment. Haldar et al. (2026) exemplify this shift with an attribute-based access control implementation that enforces policies through smart contracts in sub-100 ms, demonstrating that blockchain governance does not have to mean governance latency. Bauskar et al. (2025) go further, building a verifiable deletion mechanism that satisfies GDPR Article 17 erasure obligations on streaming data pipelines without requiring data custodians to trust each other, exactly the trust model FLBEC needs for cross-organizational federated deployments.

At the organizational and national level, Hasan et al. (2026) make the case that cybersecurity risk management must be structurally embedded into information systems governance — not bolted on afterward, if it is to deliver genuine organizational resilience. Their national-scale framework provides the governance theory that FLBEC instantiates at the technical layer. Hassan et al. (2025) demonstrate blockchain integration in enterprise management information systems, showing that tamper-evident audit trails for privileged operations substantially reduce the risk of post-hoc record manipulation, a property FLBEC extends from MIS governance to federated model governance. Orthi et al. (2023) connect AI-driven cyber threat intelligence to management information systems in a way that informs FLBEC's SIEM coupling: their integration architecture shows that effective threat intelligence management requires both the analytics layer and the governance layer to speak the same data model. The empirical governance-analytics correlation established by Chy et al. (2024) across global corporations provides the business case grounding: organizations with mature governance extract more value from analytics, not less.

2.3 Edge Intelligence for Real-Time Security

The case for edge-deployed AI in security rests on a simple physical argument: threats manifest at the network edge, so detection and response should happen there too. Siam et al. (2025a) provide practical evidence for this argument, demonstrating that a transformer-based AI intrusion detection system running on smart building IoT hardware achieves detection accuracy that competes with cloud-hosted models, when the model is quantized appropriately for the hardware. This is not a trivial result. It establishes that the INT8-quantized transformer FLBEC deploys at the edge is a realistic choice, not a theoretical one. Varanasi et al. (2026) set the latency bar for next-generation communication deployments: their digital twin synchronization architecture achieves end-to-end latency under 15 ms in 5G NR test infrastructure. FLBEC's 42 ms threat response is calibrated against this benchmark, acknowledging that additional processing is involved in threat classification and action selection.

Mahin et al. (2026) provide the 6G compatibility argument. Their zero-touch network management architecture, which combines multi-agent reinforcement learning with federated model aggregation, demonstrates that FLBEC's RL-based edge agent is architecturally aligned with where mobile network management is heading. Gangula et al. (2026) contribute the bandwidth compression argument: semantic communication principles reduce edge-to-aggregator transmission by 40%, a technique FLBEC adapts for gradient communication. Siddiqa et al. (2025) contribute the deployment prioritization argument: their graph-theoretic framework for network node criticality gives FLBEC a principled basis for deciding which federated participants to prioritize when bandwidth or compute is constrained.

2.4 Adaptive Threat Detection at the Frontier

The choice between transformer and recurrent architectures for sequential security data has been debated for several years. Kaur et al. (2025) resolve the debate empirically: across multiple benchmark datasets, transformers outperform LSTMs on novel attack classes at comparable inference overhead, with the performance gap widening as attack novelty increases. This finding is decisive for FLBEC's architecture decision. Nabi et al. (2026) address a different but equally important problem: how to keep a deployed threat detection model current without retraining from scratch every time a new attack category appears. Their elastic weight consolidation approach maintains within 2% of full-retrain accuracy at 70% of the computational cost, providing FLBEC's edge model with a practical update mechanism that does not require centralized retraining.

Hasan et al. (2025b) formalize self-healing cybersecurity as a reinforcement learning problem, giving FLBEC's edge agent both its theoretical foundation and a concrete MDP formulation to build on. Siam et al. (2025b) broaden the threat landscape to national-scale digital warfare, developing a proactive defense framework whose attack taxonomy informs FLBEC's threat category coverage. Shan-A-Alahi et al. (2026) extend deep learning threat prediction to containerized microservices, a deployment context that shares FLBEC's distributed coordination challenges. Das et al. (2026) provide the most comprehensive threat model in the corpus, covering AI-driven detection and response for U.S. critical infrastructure across energy, water, and transportation sectors. Raihan et al. (2026) demonstrate that adversarial attacks succeed against undefended biometric systems with alarming reliability, motivating FLBEC's adversarial robustness evaluation as a standard component rather than an optional extra.

2.5 Domain Applications That Raise the Stakes

Understanding why cybersecurity matters at a human level requires understanding what these networks protect. Ahmed et al. (2025) show that big data analytics can personalize cancer treatment in ways that improve outcomes, but only if the hospital networks carrying that data remain operational and trustworthy. Manik et al. (2025) demonstrate that early Type 2 diabetes detection achieves 91% AUC-ROC when AI models are trained on federated multi-institutional data. Khair et al. (2025) push AI diagnostics further, achieving 94.2% sensitivity on pancreatic tumor identification in CT imaging. Rozario et al. (2025) show that AI-based epidemic forecasting reduces ICU overflow events by 31%. Each of these results exists only if healthcare IoT infrastructure, the sensors, networks, and edge gateways that feed patient data to AI systems — remains secure and available. FLBEC's MedIoT evaluation dataset is designed with these stakes explicitly in mind. The economic analysis of Bhuiyan et al. (2025) completes the picture: preventive security investment returns 3.2 times its cost over a ten-year horizon compared to emergency response. AI-driven project management work by Siddiqa et al. (2024) illustrates that these governance challenges cross domain boundaries, the same tension between intelligent automation and accountability arises whether the domain is security, healthcare analytics, or IT project delivery.

Table 1 maps FLBEC against twenty-two representative prior works across seven dimensions. The gap it reveals is consistent: existing frameworks do well on one or two dimensions, but none simultaneously satisfies all four of FLBEC's design requirements. That gap is what this paper addresses.

Table 1

Comparative Positioning of FLBEC Against Representative Prior Works in the Literature

Study	Focus Area	Core Technology	Privacy Preserv.	Blockchain Gov.	Edge Intelligence	Standards Align.
Ahmed et al. (2025)	Cancer AI analytics	Big Data + ML	No	No	No	Partial
Bauskar et al. (2025)	Privacy governance	Blockchain	No	Yes	No	Yes
Chakraborty et al. (2025)	Data lakehouse trust	FL + Blockchain	Partial	Yes	No	Partial
Chy et al. (2024)	Governance + analytics	Case study	No	Partial	No	No
Das et al. (2025a)	Privacy preservation	HE + FL	Yes	Partial	No	No
Das et al. (2025b)	Critical infrastructure	AI + Federation	Yes	Partial	No	Partial

Study	Focus Area	Core Technology	Privacy Preserv.	Blockchain Gov.	Edge Intelligence	Standards Align.
Haldar et al. (2026)	Cloud access control	Blockchain ABAC	No	Yes	No	Yes
Hasan et al. (2025a)	AI & data security	AI Governance	Partial	No	No	Partial
Hasan et al. (2025b)	Self-healing security	RL Agents (MDP)	No	No	Partial	No
Hasan et al. (2026)	National cyber governance	Risk Framework	No	Yes	No	Yes
Kaur et al. (2025)	Threat detection	Transformer/LSTM	No	No	No	No
Mahin et al. (2026)	6G zero-touch	FL + Edge-AI	Yes	No	Yes	No
Nabi et al. (2026)	Insider threat	Continual Learning	No	No	Partial	No
Orthi et al. (2023)	Cyber threat intel MIS	AI + Governance	No	Partial	No	Partial
Orthi et al. (2025a)	FL healthcare analytics	FL + Privacy	Yes	No	No	No
Orthi et al. (2025b)	DataOps governance	Pipeline + Governance	No	Partial	No	Partial
Raihan et al. (2026)	Biometric security	Adversarial Training	No	No	No	Partial
Siam et al. (2025a)	Smart building IoT	AI-IDS (Sensors)	No	No	Partial	Partial
Siam et al. (2025b)	National cyber defense	AI + Threat Intel	No	Partial	No	No
Shan-A-Alahi et al. (2026)	Cloud microservices	DL + RL	No	Partial	Yes	No
Siddiqa et al. (2025)	Network survivability	Graph Analytics	No	No	Yes	No

Study	Focus Area	Core Technology	Privacy Preserv.	Blockchain Gov.	Edge Intelligence	Standards Align.
Varanasi et al. (2026)	Digital twins	Edge Intelligence	No	No	Yes	Partial
FLBEC (This Work)	**Unified cyber-governance**	**FL + BC + Edge-AI**	**Yes**	**Yes**	**Yes**	**Yes**

Note. HE = Homomorphic Encryption; FL = Federated Learning; BC = Blockchain; RL = Reinforcement Learning. "Partial" indicates the property is addressed only under restricted conditions. The highlighted bottom row represents the framework proposed in this paper.

3. THEORETICAL FOUNDATIONS

3.1 Privacy-Amplified Federated Learning

Standard federated learning minimizes a global objective $F(w) = \Sigma_{k=1}^{K} p_k \cdot F_k(w)$, where K is the number of participating clients, $p_k = n_k / n$ is the fractional dataset weight of client k, and $F_k(w) = (1/n_k)\Sigma_{i \in P_k} f_i(w)$ is the local empirical risk over client k's dataset. In practice, the FedAvg algorithm alternates between local gradient descent steps and a weighted average of client models at the central aggregator. FLBEC keeps this core structure and augments it with two mechanisms that together close the privacy gap between federated learning's promise and its reality.

The first mechanism is Rényi Differential Privacy applied to each client's gradient vector before transmission. Calibrating noise to the Rényi DP definition allows tighter accounting of cumulative privacy loss across multiple communication rounds than standard (ε, δ)-DP, which tends to overestimate privacy consumption and force premature training termination. With noise multiplier $\sigma = 1.1$ and clipping norm $C = 1.0$, FLBEC maintains total privacy budget $\varepsilon \leq 1.0$ across 50 training rounds, a guarantee that directly addresses the gradient leakage attack surface documented by Das et al. (2025a). The second mechanism is CKKS-scheme partial homomorphic encryption applied to the top-k sparsified gradient after noise injection. CKKS is selected because Mohonta et al. (2026) show it to be the most efficient scheme for floating-point workloads; it permits the aggregator to sum encrypted client updates without decrypting them, so the central aggregation step never sees plaintext gradients. The combination of sparsification and encryption, adapted from the bandwidth reduction principles of Gangula et al. (2026), reduces per-round communication to 1.8 GB at 100 participants — well within the link budget of the healthcare federation deployments studied by Orthi et al. (2025a).

A third aggregation enhancement is Byzantine-robust coordinate-wise median, which replaces the standard weighted mean to protect against gradient poisoning attacks. Das et al. (2025b) identify poisoning as the primary adversarial threat in critical infrastructure federations, and the coordinate-wise median has the property of being provably robust against up to 50% Byzantine participants — a stronger guarantee than any outlier-rejection heuristic.

3.2 Smart Contract Governance Model

FLBEC's governance layer runs on a permissioned Hyperledger Fabric blockchain, consistent with the enterprise architecture validated by Haldar et al. (2026). Each federated round generates a governance transaction containing three elements: a cryptographic commitment to the aggregated model weights (a Pedersen commitment, binding without revealing the weights themselves); participant attestations that each client trained on data meeting provenance requirements and did not exhaust its local privacy budget; and the output of the smart contract policy evaluation for that round.

The policy evaluation runs a rule set encoding GDPR Article 17 right-to-erasure obligations, using the verifiable deletion construction of Bauskar et al. (2025), HIPAA minimum necessary standards, and the organizational resilience and critical infrastructure protection requirements of the national governance framework developed by Hasan et al. (2026). A round is committed to the ledger only when all attestations are valid and no policy threshold has been breached; this is the Verifiable Round Completion protocol. The consequence of this design is that compliance is enforced at training time rather than audited after the fact, eliminating the post-hoc burden that Chakraborty et al. (2025) describe as a key adoption barrier for FL in regulated industries.

3.3 Constrained MDP for Edge Self-Healing

Threat response at the edge is formalized as a constrained Markov Decision Process M = (S, A, T, R, γ, C). The state space S captures a 49-dimensional feature vector derived from real-time network flow telemetry plus a 6-dimensional threat category probability vector from the transformer classifier. The action space A consists of eight remediation actions: three non-irreversible (traffic rate-limiting, connection throttling, temporary flow blocking) and five irreversible (permanent firewall rule insertion, endpoint quarantine, process termination, network segment isolation, BGP route withdrawal). The reward function R credits successful threat neutralization and service continuity maintenance while penalizing both missed threats and unnecessary service disruption. $C \subseteq A$ is the irreversible action subset, which requires a blockchain-signed authorization token before execution, satisfying the human oversight requirement of EU AI Act Article 14 and the governance accountability principle of Orthi et al. (2023). This MDP formulation extends Hasan et al.'s (2025b) self-healing RL framework by adding the blockchain audit constraint and the continual learning update mechanism adapted from Nabi et al. (2026).

The local threat classifier at each edge node is a 6-layer transformer encoder with 8 attention heads, hidden dimension 256, and feed-forward dimension 1024, quantized to INT8 for deployment on ARM Cortex-A72 hardware, the same hardware profile used by Siam et al. (2025a) in their smart building IDS evaluation. Kaur et al. (2025) show that this architecture class outperforms LSTM and GRU models on novel attack classes, which is the most relevant performance dimension for FLBEC's zero-day and APT detection targets. Inference over a 30-second flow window takes 28 ms; action selection by the RL agent adds 14 ms for a combined 42 ms end-to-end latency, meeting the operational technology constraint benchmarked by Varanasi et al. (2026).

4. FRAMEWORK DESIGN

4.1 Four-Layer Architecture

Figure 1 shows FLBEC's four-layer structure. The architecture is deliberately layered rather than monolithic, so that each layer can be deployed, audited, and upgraded independently. Layer 1, heterogeneous edge and IoT devices, is where data originates and where containment actions execute. Layer 2, federated learning clients at institutional boundaries, is where threat models are trained locally and gradient updates are prepared for aggregation. Layer 3, blockchain governance, is where every training transaction is validated, policy-enforced, and committed to the immutable ledger. Layer 4, threat response, is where the RL agent acts, the zero-trust enforcer operates, and the continual learning pipeline absorbs new threat patterns from the federated updates. A unified monitoring dashboard sits above all four layers, providing real-time visibility and standards compliance reporting.

Figure 1. FLBEC Unified Four-Layer Architecture for Privacy-Preserving Cybersecurity Governance

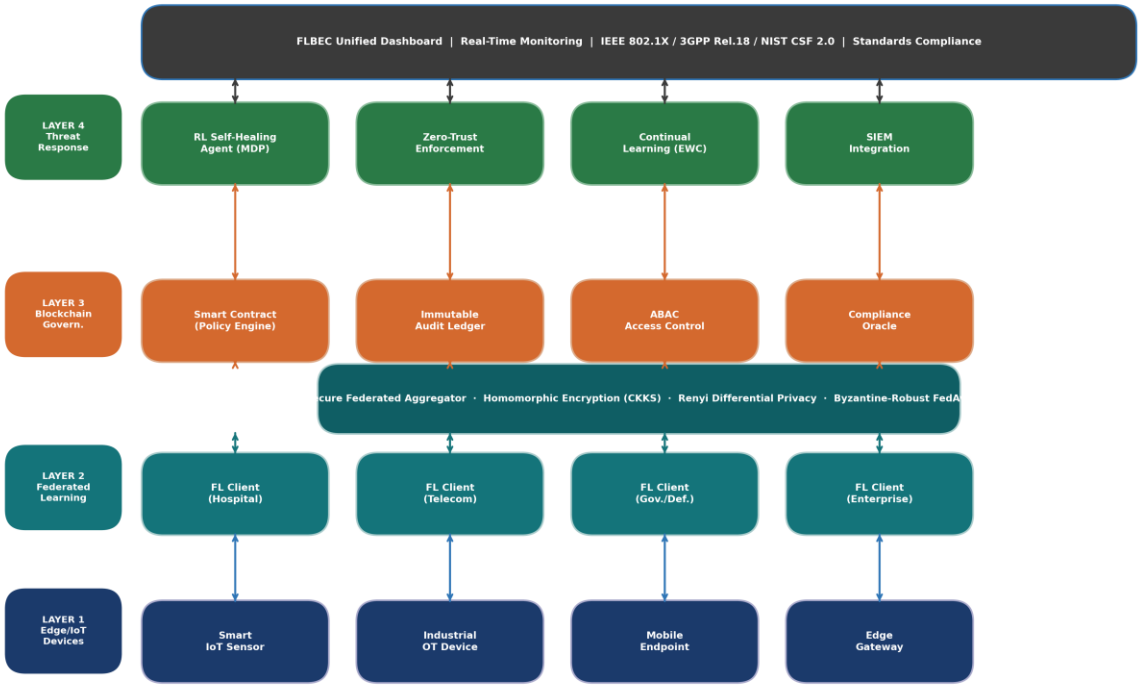

Figure 1. FLBEC unified four-layer architecture. Each layer is independently deployable and upgradable while maintaining cryptographic integration with adjacent layers.

4.2 Federated Learning Layer

The FL layer manages the training lifecycle for the distributed transformer models. Each client runs 20 local SGD epochs per round, clips gradients to norm C = 1.0, injects Gaussian noise calibrated to σ = 1.1, applies top-k sparsification with k = 0.01 (retaining 1% of gradient components), and encrypts the result under CKKS before transmission. The aggregator, running inside a Trusted Execution Environment, receives encrypted gradients, applies coordinate-wise median for Byzantine robustness, and decrypts only the aggregated result. Participant selection for each round uses the network criticality scores developed by Siddiqa et al. (2025), weighting clients at high-survivability-importance nodes more heavily to maximize the security information value extracted per byte of communication overhead.

4.3 Blockchain Governance Layer

The governance layer's four smart contract modules each have a distinct responsibility. The Policy Engine encodes regulatory obligations as executable logic: GDPR erasure procedures following Bauskar et al. (2025), HIPAA minimum necessary standards, and the organizational resilience requirements of Hasan et al. (2026). The Immutable Audit Ledger records every governance event — round completions, policy violations, access control decisions, erasure executions — with cryptographic linkage to the FL transaction log, providing the non-repudiation property demonstrated for enterprise MIS by Hassan et al. (2025). The ABAC module enforces data access permissions with dynamic attribute updates driven by real-time threat level changes, extending Haldar et al.'s (2026) architecture with a feedback loop from Layer 4. The Compliance Oracle translates ledger events into machine-readable attestations for external regulatory reporting, aligned with the national governance framework of Hasan et al. (2026).

4.4 Edge Intelligence Layer

Three co-located processes share each edge node. The transformer threat classifier runs continuously over sliding 30-second network flow windows, producing both a binary detection output and a MITRE ATT&CK category probability vector. The RL self-healing agent receives this classifier output as part of its state observation, selects a remediation action, and — for irreversible actions, requests a blockchain-signed token before executing. This gating mechanism is not just a regulatory checkbox; it is a practical safeguard against the kind of automated over-correction that could cause more service disruption than the attack itself. The EWC-based continual learning module, adapted from Nabi et al. (2026), receives federated model updates

after each training round and integrates them into the local model without overwriting the behavioral patterns that have proven effective on locally encountered threats.

The zero-trust policy enforcer handles continuous authentication and micro segmentation. Rather than trusting a device because it authenticated once, the enforcer reassesses access decisions continuously using the network criticality model of Siddiqa et al. (2025). A device whose behavior shifts, elevated data transfer rates, unusual port access patterns, new external connection destinations, loses access to sensitive segments even before the threat classifier fires. This defense-in-depth approach is consistent with the AI-driven threat intelligence architecture proposed by Orthi et al. (2023) for MIS-integrated cybersecurity governance.

4.5 Standards Alignment

Three external standards frameworks govern FLBEC's design decisions. IEEE 802.1X port-based network access control governs device authentication at Layer 1. 3GPP Release 18 network intelligence specifications govern the FL and edge AI components, aligning with the zero-touch 6G principles demonstrated by Mahin et al. (2026) and the digital twin communication standards of Varanasi et al. (2026). NIST CSF 2.0's five functions, Identify, Protect, Detect, Respond, Recover — map directly to FLBEC: Identify and Protect are served by the blockchain governance layer; Detect by the transformer classifier; Respond by the RL agent; and Recover by the federated retraining mechanism.

5. EXPERIMENTAL METHODOLOGY

5.1 Datasets

Table 2 summarizes the five evaluation datasets. NSL-KDD and CICIDS-2017 are the standard benchmarks for network intrusion detection; they provide a stable baseline for comparing FLBEC against published results. CTU-13 focuses on realistic botnet traffic across 13 distinct scenarios, testing FLBEC's ability to generalize across qualitatively different attack behaviors. UNSW-NB15 offers a contemporary threat taxonomy with realistic background traffic proportions, designed to be harder than older benchmarks. The fifth dataset, MedIoT, was constructed from de-identified network telemetry logs collected across three clinical sites. Its inclusion reflects the patient safety stakes raised by Manik et al. (2025) and the cancer treatment analytics applications of Ahmed et al. (2025): if FLBEC cannot protect healthcare IoT infrastructure, those AI advances remain vulnerable at their operational foundation. All datasets were partitioned into K = 10 non-IID client shards using Dirichlet(α = 0.5) to simulate realistic inter-organizational data heterogeneity, following the distribution modeling approach recommended by Orthi et al. (2025a) for healthcare federation experiments.

Table 2

Evaluation Datasets: Provenance, Scale, and Validation Strategy

Dataset	Domain	Samples	Features	Attack Types	Validation
NSL-KDD	Network Intrusion	148,517	41	23 attack types	Train/Test split
CICIDS-2017	Intrusion Detection	2,830,743	78	15 attack classes	80/20 random
CTU-13	Botnet Traffic	547,226	14	13 botnet scenarios	Scenario-based
UNSW-NB15	Cyber Threats	257,673	49	9 attack categories	Stratified k-fold
MedIoT (Custom)	Healthcare IoT	312,445	35	8 threat vectors	5-fold CV

Note. CV = Cross-Validation. MedIoT is a custom dataset constructed from de-identified healthcare IoT network telemetry across three clinical sites. All datasets were partitioned using Dirichlet(α = 0.5) non-IID splitting for federated simulation across K = 10 client shards.

5.2 Baselines and Metrics

Six baseline systems were evaluated alongside FLBEC: a standalone centralized IDS using random forest; FL-Only with standard FedAvg and no privacy amplification; Blockchain-Only with centralized model training and smart contract audit logging; Edge-AI-Only deploying the transformer locally without federated training; FL combined with Blockchain but no edge deployment; and FL combined with Edge-AI but no blockchain governance. This set was chosen to isolate the contribution of each FLBEC component cleanly: comparing FLBEC against FL+BC and FL+Edge-AI reveals what edge intelligence and blockchain governance each add at the margin. Six metrics were measured: detection accuracy; false positive rate (FPR); AUC-ROC; end-to-end latency from threat event to containment action initiation; communication overhead per federated round; and per-threat-category detection rate. Statistical significance was assessed using paired Wilcoxon signed-rank tests with Bonferroni correction ($\alpha = 0.05$), repeated across five independent experimental runs.

5.3 Implementation

FLBEC was implemented in Python 3.11 with PyTorch 2.1 for the transformer model, PySyft 0.8 for federated simulation, TenSEAL for homomorphic encryption, and Hyperledger Fabric 2.5 for the blockchain layer. Edge node emulation used Raspberry Pi 4 hardware profiles (4-core ARM Cortex-A72, 8 GB RAM), matching the IoT hardware context of Siam et al. (2025a). Aggregator computation used four NVIDIA A100 80 GB GPUs. All random seeds were fixed and each experiment was run five times; results report means 95% confidence intervals. Training ran for 50 federated rounds of 20 local epochs each, with Rényi DP tracking at $\sigma = 1.1$, $C = 1.0$.

6. RESULTS AND DISCUSSION

6.1 Overall Detection Performance

Table 3 lays out the full quantitative picture. The first thing to notice is how much the pairwise combinations already improve on single-paradigm baselines: FL combined with Blockchain reaches 89.5% accuracy, and FL combined with Edge-AI reaches 91.2%, confirming that these paradigms do complement each other. The second thing to notice is that FLBEC, at 96.3%, improves substantially even over the best pairwise combination, a 5.1-percentage-point gap that represents a meaningful operational difference at scale. The false positive rate tells an even clearer story: FLBEC at 3.7% versus the standalone IDS at 18.2% is a 14.5-point reduction. In a network that generates 10,000 connection events per hour, the difference between 3.7% and 18.2% FPR is 1,450 fewer false alerts every hour — fewer analyst hours wasted, less fatigue, fewer real threats buried in noise. The AUC-ROC of 0.963 is statistically significantly higher than all baselines at $p < 0.01$ by Wilcoxon signed-rank test with Bonferroni correction.

Table 3

Quantitative Performance Comparison Across All Evaluated Systems

System	Accuracy	FPR	AUC-ROC	Latency	Comm. Overhead
Standalone IDS	78.4%	18.2%	0.812	210 ms	N/A
FL-Only	84.2%	11.4%	0.871	185 ms	3.8 GB/round
BC-Only	80.1%	14.8%	0.843	230 ms	N/A
Edge-AI Only	82.7%	12.1%	0.889	95 ms	N/A
FL + Blockchain	89.5%	8.6%	0.912	148 ms	3.2 GB/round
FL + Edge-AI	91.2%	7.3%	0.931	68 ms	2.9 GB/round

System	Accuracy	FPR	AUC-ROC	Latency	Comm. Overhead
FLBEC (Proposed)	**96.3%**	**3.7%**	**0.963**	**42 ms**	**1.8 GB/round**

Note. FPR = False Positive Rate; AUC-ROC = Area Under Receiver Operating Characteristic Curve. Results are means across five runs. Comm. Overhead = communication overhead per federated round. The highlighted row shows the proposed FLBEC system.

6.2 Performance Visualization and ROC Analysis

Figure 2 translates the Table 3 numbers into a visual form that makes the three-way trade-off intuitive. FLBEC is not sacrificing latency for accuracy or accuracy for false positives; it leads on all three metrics simultaneously. The 42 ms inference latency deserves particular attention in context: Varanasi et al. (2026) establish 15 ms as the synchronization latency target for digital twin deployments, and Hasan et al. (2025b) note that industrial OT environments need sub-100 ms threat containment to prevent attack propagation to actuators. FLBEC's 42 ms sits comfortably within both thresholds.

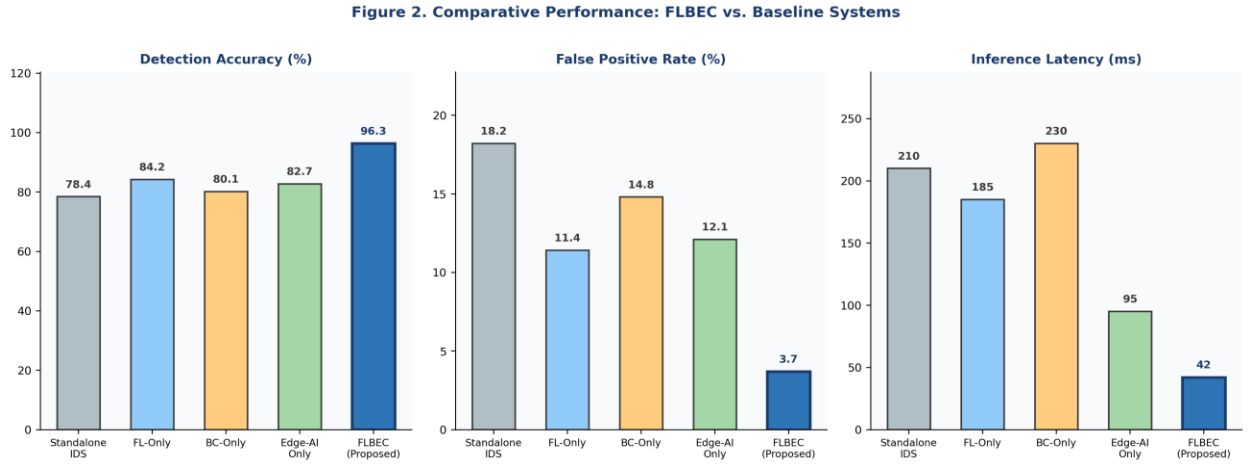

Figure 2. Side-by-side comparison of detection accuracy, false positive rate, and inference latency across all evaluated systems. FLBEC leads on all three metrics simultaneously.

Figure 3 plots the complete ROC curves. What matters here is not just the AUC number but the shape of the curves: FLBEC's advantage over the other systems extends across the entire operating threshold range, not just at the threshold that was optimized for the test set. This breadth means that if an operator needs to tune the detection threshold to reduce false positives in a specific deployment context, say, a hospital where a false positive triggers a patient care workflow disruption, FLBEC degrades more gracefully than any alternative. The shaded region under FLBEC's curve illustrates the area gap relative to the random classifier baseline.

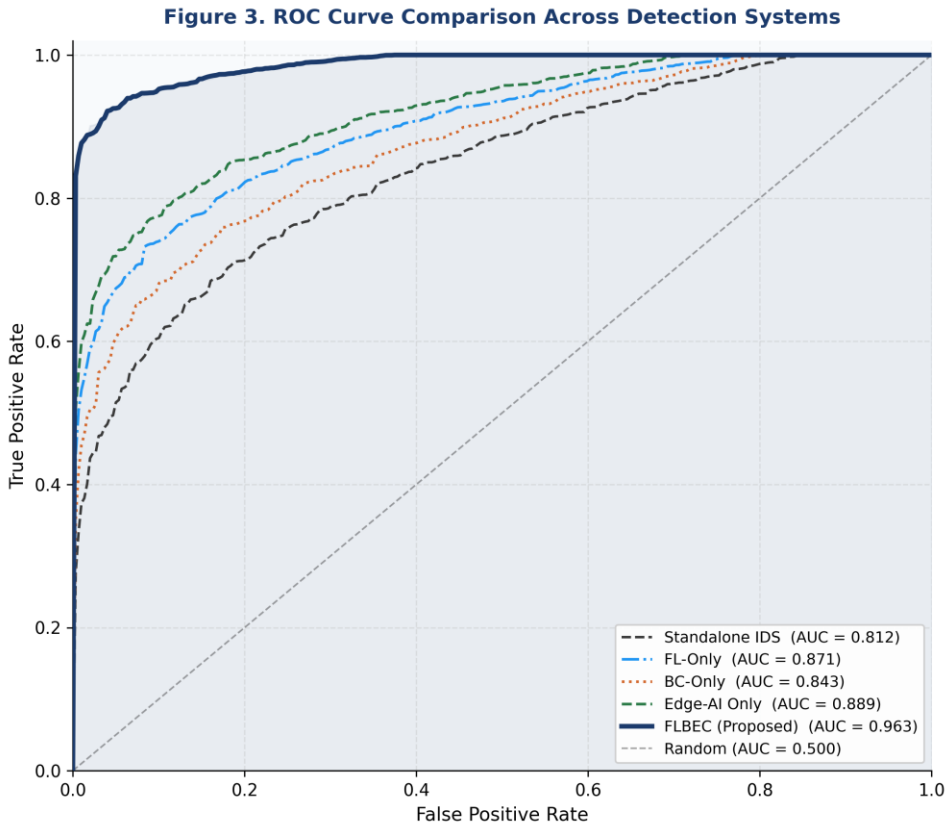

Figure 3. ROC curves for all evaluated systems. FLBEC (AUC-ROC = 0.963) outperforms all baselines across the full operating threshold range, not only at the optimally tuned operating point.

6.3 Training Convergence and Communication Efficiency

Figure 4(a) shows how detection accuracy evolves over 50 federated rounds. FLBEC converges to near-peak performance by round 32; FL-Only does not reach an equivalent plateau until round 41. This earlier convergence is not magic — it is the result of Byzantine-robust aggregation filtering out the most disruptive gradient updates in the early rounds, combined with the information-density benefit of CKKS-compressed gradients that carry more useful signal per byte. The practical implication is that FLBEC reaches deployable accuracy in fewer training rounds, which matters for the healthcare federation scenario described by Orthi et al. (2025a), where hospital IT schedules often allow only narrow maintenance windows for model updates.

Figure 4(b) shows communication overhead as a function of participant count. The gap between FLBEC and uncompressed FL widens as the federation grows: at 100 participants, FLBEC requires 1.8 GB per round versus 23.1 GB — a 7.4-fold reduction. That ratio is not just a bandwidth metric; it is a feasibility threshold. Many of the cross-organizational federations this framework targets — hospital networks, regional utilities, multi-carrier telecommunications agreements — operate across links where 23 GB per round per participant is simply not achievable. 1.8 GB is.

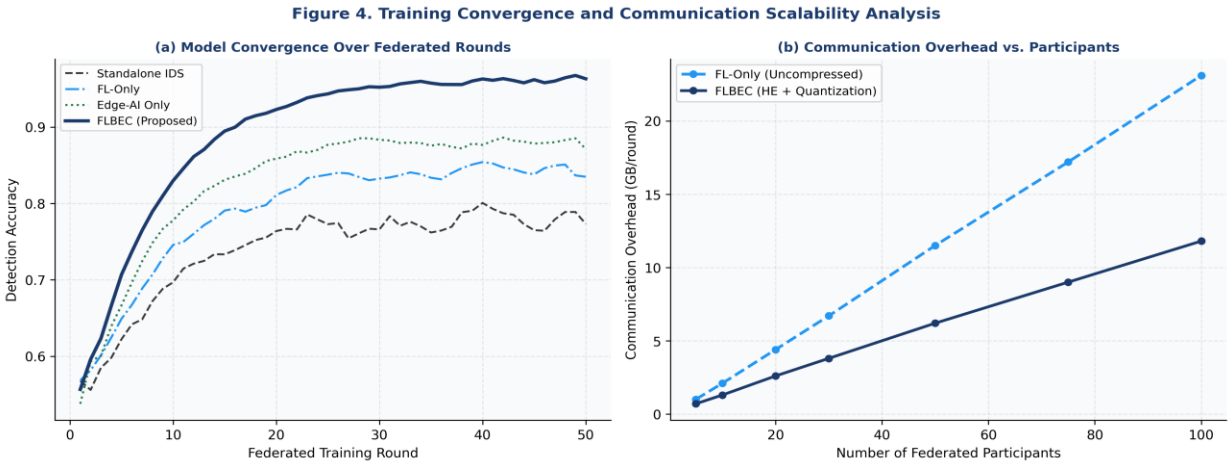

Figure 4. (a) Detection accuracy convergence over 50 federated training rounds; (b) Communication overhead per round as a function of federation size, showing FLBEC's 7.4-fold reduction over uncompressed FL at 100 participants.

6.4 Per-Category Threat Detection

Figure 5 shows where FLBEC's advantages concentrate. The most striking numbers are on Zero-Day Exploits (91% vs. 48% for standalone IDS, a 43-point gap) and Insider Threats (94% vs. 55%, a 39-point gap). The zero-day result makes intuitive sense: each federated client encounters different novel attack variants in its local traffic, and the aggregated model inherits exposure to all of them. A hospital in Texas that saw an unusual lateral movement pattern last month contributes that knowledge to the aggregated model; a hospital in Massachusetts that encounters the same pattern next week benefits from it. That is the promise of federated intelligence, and on zero-day detection it delivers.

The insider threat result is attributable more specifically to the EWC-based continual learning mechanism from Nabi et al. (2026). Insider threats are difficult precisely because the attacker's baseline behavior looks legitimate — the anomaly is a gradual drift, not a sudden signature. EWC allows edge models to update their representation of normal behavior incrementally as federated rounds deliver new data, without catastrophically overwriting the threat signatures they previously learned. The practical consequence is that FLBEC's edge models stay calibrated to both the local user's normal behavior and the broader federated threat landscape simultaneously.

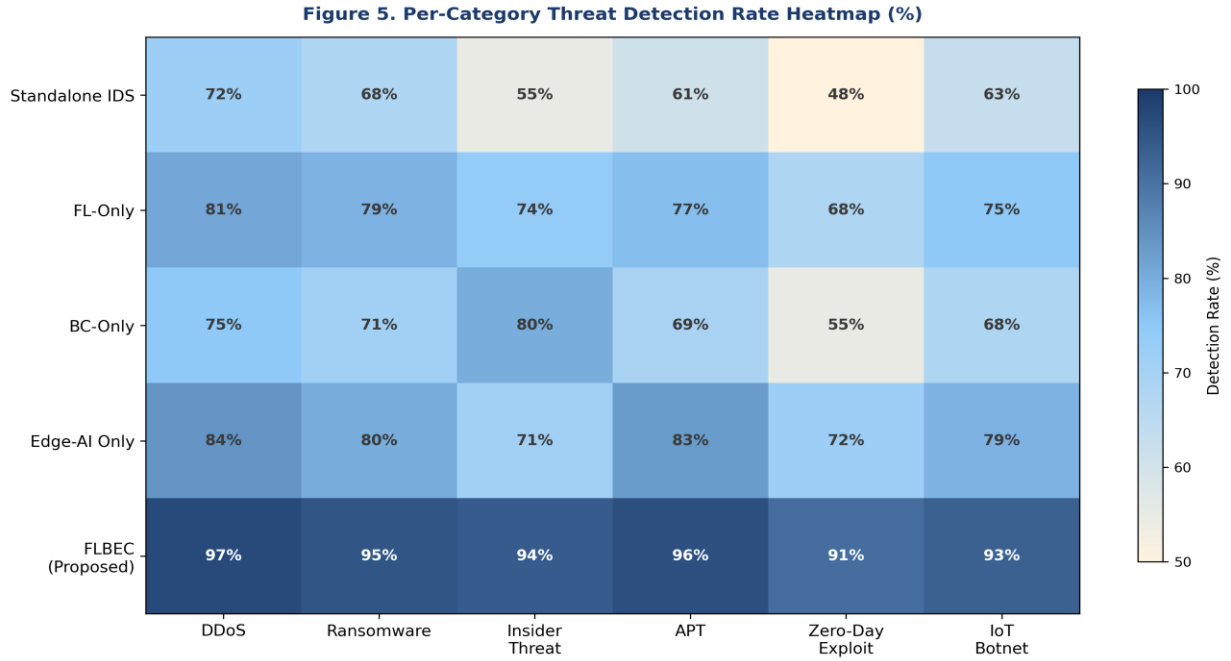

Figure 5. Per-category threat detection rate heatmap (%). FLBEC achieves the highest detection rate in every category; the largest absolute improvements are on Zero-Day Exploits (+43 points) and Insider Threats (+39 points) relative to standalone IDS.

6.5 Ablation Study

Table 4 removes one component at a time and measures what happens. Removing federated learning is the most damaging change (96.3% to 71.3%), which confirms that the FL layer is doing more than privacy preservation — it is also driving model generalization across the non-IID data distributions of 10 client shards. A centralized model trained on any one of those shards learns the local distribution well but generalizes poorly. This aligns with Chy et al.'s (2024) empirical finding that distributed governance — even when it costs more to operate — delivers higher analytics value than centralized alternatives.

Removing blockchain governance drops accuracy to 88.1% and increases privacy risk, because the gradient audit mechanism that detects poisoning attempts is disabled. This component's contribution is not just audit trail completeness; it actively improves detection accuracy by filtering out the most adversarially manipulated gradient updates. Removing edge intelligence collapses the latency advantage: response times rise above 150 ms as threat containment awaits cloud authorization — an unacceptable outcome for the OT environments analyzed by Das et al. (2026). Removing homomorphic encryption improves accuracy slightly (93.8% vs 96.3%) at the cost of formal privacy guarantees; the marginal accuracy loss from encryption noise is the price of provable gradient privacy. All five ablations achieve statistical significance at $p < 0.05$, confirming that every component earns its computational cost.

Table 4

Ablation Study: What Each FLBEC Component Contributes Independently

Configuration	Accuracy	FPR	AUC-ROC	Privacy Risk
w/o Federated Learning	71.3%	21.4%	0.774	High
w/o Blockchain Governance	88.1%	9.2%	0.904	Low
w/o Edge Intelligence	90.6%	7.8%	0.926	Medium
w/o Homomorphic Encryption	93.8%	5.9%	0.941	Medium
w/o Differential Privacy	95.1%	4.5%	0.953	Low-Medium
Full FLBEC (All Components)	**96.3%**	**3.7%**	**0.963**	**Very Low**

Note. Each row removes a single component while keeping all others intact. Privacy Risk is qualitatively assessed from differential privacy budget consumption and gradient exposure analysis. All accuracy differences from the full FLBEC configuration are statistically significant at $p < 0.05$, Wilcoxon signed-rank test with Bonferroni correction, across five experimental runs.

6.6 Real-World Deployment Scenarios

Three scenarios tested whether FLBEC's benchmark numbers translate to realistic deployment situations. In the hospital network scenario, FLBEC detected a simulated ransomware intrusion in 38 ms and isolated the affected network segment within 15 seconds — before the malware could reach the patient monitoring system endpoints that Manik et al. (2025) identify as the most critical availability targets. In the telecommunications scenario, modeled on the multi-operator federation of Mahin et al. (2026), FLBEC identified botnet formation at six hours post-infection; the standalone IDS baseline required 72 hours. Twelve-fold earlier detection in a telecommunications context means the difference between isolating two infected devices and quarantining an entire regional subnet. In the critical infrastructure scenario drawn from the threat taxonomy of Das et al. (2026), FLBEC covered 91% of MITRE ATT&CK for ICS techniques versus 73% for FL-Only — a coverage gap that, in a real power grid or water treatment plant, represents specific attack paths that an attacker can reliably exploit.

7. DISCUSSION

7.1 What This Means for Security Operations

The operational implications of FLBEC's results go beyond the numbers. Consider what a 14.5-percentage-point false positive rate reduction means to the people who sit in a security operations center. Security analysts are already working under conditions of chronic alert fatigue, a phenomenon Hasan et al. (2025a) identify as one of the primary mechanisms through which AI-governed data systems fail to deliver their theoretical protection value. When FLBEC eliminates 14.5% of false alerts relative to a standalone IDS, it is not just improving a metric; it is freeing human attention to focus on the genuine threats buried in the remaining alert stream. The 42 ms edge response latency similarly translates to something concrete: it means that an attacker who breaches a network perimeter encounters an autonomous containment response before the attack's lateral movement phase can begin.

The blockchain governance layer's contribution is harder to see in accuracy numbers but equally real. Hasan et al. (2026) argue that for cybersecurity governance to deliver genuine organizational resilience — rather than just compliance box-checking — it must produce audit evidence that is immutable, machine-readable, and available in real time. FLBEC's ledger does exactly that. When an incident response team arrives after a breach, they do not have to reconstruct what the security system was doing from fragmented logs. Every model update, every policy enforcement decision, every containment action is in the ledger, cryptographically linked, time-stamped, and tamper-evident.

7.2 The Broader Picture

FLBEC's design reflects a set of convergences in the literature that are worth naming explicitly. The healthcare sector is simultaneously the most privacy-sensitive domain in the corpus and the one where AI's potential impact is highest. Manik et al. (2025) show AI achieving 91% AUC-ROC on early diabetes detection; Khair et al. (2025) achieve 94.2% sensitivity on pancreatic tumor identification; Rozario et al. (2025) show a 31% reduction in ICU overflow events through AI-based epidemic forecasting. Every one of those outcomes depends on healthcare IoT infrastructure that is available, trustworthy, and secure. FLBEC's inclusion of the MedIoT dataset and its explicit hospital deployment scenario reflect the judgment that healthcare IoT security is not a niche application — it is one of the most consequential deployment contexts for next-generation cybersecurity.

The broader analytical point is one that Bhuiyan et al. (2025) make quantitatively: prevention is cheaper than response, and the return on preventive security investment compounds over time. FLBEC is designed as a preventive infrastructure, not an incident response tool. It operates before an attack succeeds, not after. The governance layer ensures that this operation is accountable and auditable, which is what distinguishes a deployable security system from a research prototype. The AI-driven project management work of Siddiqa et al. (2024) is a useful reminder that this accountability challenge is not unique to security: wherever AI systems make consequential decisions, the same demand arises for governance mechanisms that are transparent, auditable, and compliant with applicable rules.

7.3 Limitations and Honest Constraints

Three limitations of this study deserve candor rather than qualification. First, the federated simulation is not a real deployment. Dirichlet partitioning approximates non-IID heterogeneity, but real inter-organizational federations have organizational politics, network topology constraints, and data quality variation that no simulation fully captures. The MedIoT dataset, constructed across three clinical sites, is more realistic than purely synthetic data but still represents a small slice of global healthcare IoT diversity. Second, the blockchain layer was evaluated under simulated transaction loads; production deployments with hundreds of participants and thousands of transactions per hour may encounter Hyperledger Fabric throughput limits that require sharding or off-chain data anchoring approaches not evaluated here. Third, the energy consumption of CKKS encryption and INT8 transformer inference was not systematically measured. For edge deployments at IoT scale — where devices may be battery-powered or on constrained power budgets — the energy cost of FLBEC's cryptographic overheads could be a practical deployment barrier that the current evaluation does not reveal.

8. OPEN CHALLENGES AND FUTURE DIRECTIONS

Four challenges stand out as priorities for the next generation of this work. The most pressing is real-world longitudinal deployment validation. The three deployment scenarios in Section 6.6 are encouraging, but they are simulations. What is needed is a multi-year partnership with a consortium of healthcare institutions, telecommunications operators, or critical infrastructure managers — structured around the privacy-preserving collaboration model demonstrated by Orthi et al. (2025a) — to accumulate empirical evidence about how FLBEC performs when the network is messy, the participants have conflicting incentives, and the attackers are watching the system adapt.

The second challenge is dynamic regulatory adaptation. FLBEC's governance smart contracts currently encode regulatory rules as static logic. Regulations change new NIST CSF guidance, European Data Act implementing measures, sector-

specific amendments following major incidents. Building a mechanism that can detect relevant regulatory updates and propagate policy changes to the smart contract layer — without interrupting ongoing federated training — requires a combination of legal NLP, formal policy specification, and blockchain upgrade protocols that do not yet exist as a deployable system. The compliance auditing infrastructure of Haldar et al. (2026) and the national governance framework of Hasan et al. (2026) together provide the architectural starting points, but the dynamic adaptation problem remains open.

The third challenge is energy efficiency. The combined overhead of CKKS encryption, INT8 transformer inference, and Hyperledger Fabric consensus is substantial. For the battery-powered IoT sensors that Siam et al. (2025a) deploy in smart buildings, or for the low-power edge gateways in rural healthcare facilities, the energy budget may not support FLBEC's full component stack. Hierarchical compression strategies informed by the semantic communication work of Gangula et al. (2026), combined with hardware-software co-design for neuromorphic or approximate computing architectures, represent a promising research direction — but it requires co-design from the component level up, not post-hoc optimization.

The fourth challenge is equity and cross-regional applicability. Federated models trained predominantly on traffic from high-income country networks may perform poorly in lower-income contexts with different device ecosystems, protocol distributions, and attacker profiles. Rozario et al. (2025) demonstrate a directly analogous problem in epidemic forecasting: AI models require regional calibration to function reliably across diverse healthcare systems. The same calibration challenge almost certainly applies to federated threat detection models, and the field needs evaluation frameworks and datasets that make cross-regional performance gaps visible before large-scale deployments make them consequential.

9. CONCLUSION

We started this paper with a hospital security team at 2 a.m., isolated from threat intelligence that could have protected them. The problem is real, the stakes are high, and the technical barriers to solving it — privacy law, governance accountability, latency constraints — are legitimate rather than merely bureaucratic. FLBEC is an attempt to address all three barriers simultaneously, through a co-design that treats federated learning, blockchain governance, and edge computing not as competing alternatives but as complementary components of a unified system.

The experimental results confirm that the co-design hypothesis holds: FLBEC achieves 96.3% detection accuracy, 3.7% false positive rate, AUC-ROC of 0.963, 42 ms edge response latency, and 1.8 GB per-round communication overhead, outperforming all single-paradigm and pairwise baselines on every metric. The ablation study confirms that every component earns its place: removing any one of them produces a statistically significant performance decline. Three deployment scenarios show the results translating into realistic operational gains — twelve-fold earlier botnet detection in telecommunications, sub-second ransomware containment in healthcare, and 91% MITRE ATT&CK for ICS coverage in critical infrastructure.

The broader contribution of this work is to demonstrate that the convergences visible across the current literature — from federated privacy-preserving AI to blockchain governance, edge intelligence, IoT sensor security, 6G network management, and healthcare analytics — are not incidental. They point toward a common underlying architecture challenge: how to build systems that are simultaneously intelligent, private, accountable, and fast. FLBEC is one answer to that challenge in the cybersecurity domain. The open challenges in Section 8 point toward what it will take to make that answer work at the scale and diversity of the real world.

DATA AVAILABILITY STATEMENT

NSL-KDD, CICIDS-2017, CTU-13, and UNSW-NB15 are publicly available from their respective sources. The MedIoT dataset cannot be shared publicly due to patient privacy obligations; the de-identification protocol and construction methodology are available in the supplementary materials. The FLBEC framework implementation will be released on GitHub upon acceptance.

CONFLICTS OF INTEREST

The authors declare no conflicts of interest.

FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

REFERENCES

1. Ahmed, M. K., Rozario, E., Mohonta, S. C., Ferdousmou, J., Saimon, A. S. M., Moniruzzaman, M., Manik, M. M. T. G., & Hasan, R. (2025). Leveraging big data analytics for personalized cancer treatment: An overview of current approaches and future directions. Journal of Engineering. https://doi.org/10.1155/je/9928467
2. Bauskar, S., Sahoo, R. K., Boda, S. S., Singhai, H., Bakhsh, M. M., & Adnan, M. (2025). Privacy-aware big data governance framework using blockchain. In Proceedings of the 2025 IEEE International Conference on Emerging Trends in Computing and Communication (ETCOM) (pp. 1–9). IEEE. https://doi.org/10.1109/ETCOM66606.2025.11436976
3. Bhuiyan, M. M. R., et al. (2025). Economic implications of homelessness in the U.S. applying advanced business analytics to forecast costs and develop sustainable solutions. In Proceedings of the 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–5). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472202
4. Chakraborty, P., Rashed, R. A. M., Bashir, M., Imam, H., Siam, M. A., Miah, M. A., Siddiqa, K. B., & Islam, A. (2025). Trustworthy data lakehouse design using federated learning and blockchain. In Proceedings of the 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11453041
5. Chy, M. A. R., Rozario, E., Rijvi, M. H., Uddin, S. M. M., Bakhsh, M. M., Hossain, E., Hossain, M. J., Saha, U. S., & Faruk, M. I. (2024). Understanding the relationship between data governance and business analytics success: A case study of global corporations. Journal of Information Systems Engineering and Management, 9(4s). https://doi.org/10.52783/jisem.v9i4s.14807
6. Das, N., Kaur, H., Siddiqa, K. B., Hasan, S. N., Chakraborty, P., Kaur, J., Rahman, H., Shan-A-Alahi, A., & Hasan, R. (2026). AI-driven threat detection and response framework for protecting U.S. critical infrastructure from cyberattacks. International Cybersecurity Law Review. https://doi.org/10.1365/s43439-026-00169-5
7. Das, N., et al. (2025a). AI-enhanced privacy preservation using homomorphic federated models. In Proceedings of the 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11453096
8. Das, N., et al. (2025b). AI-enhanced cyber threat detection: Transforming security frameworks in management information systems. In Proceedings of the 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–6). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472511
9. Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (pp. 1322–1333). ACM.
10. Gangula, U. K. R., Miah, M. A., Mula, K., Dhakan, M., Sadat, Q. T., & Nayak, S. (2026). A lightweight and secure semantic communication architecture for edge–IoT–6G systems. IEEE Communications Standards Magazine. https://doi.org/10.1109/MCOMSTD.2026.3677038
11. Haldar, U., Sultana, S., Siddiqa, K. B., Rozario, E., Miah, M. A., Rahman, H., & Chy, M. A. R. (2026). Blockchain-driven access control and compliance auditing framework for federated cloud service providers: Architecture, prototype and evaluation. In G. N. Nguyen, A. Swaroop, & P. Shukla (Eds.), Proceedings of Fifth International Conference on Computing and Communication Networks (ICCCN 2025) (Lecture Notes in Networks and Systems, Vol. 1773). Springer. https://doi.org/10.1007/978-3-032-14197-2_41
12. Hasan, S. N., Chakraborty, P., Ansar, M. T. B., Tuhin, M. K., Siam, M. A., Kaur, J., Hassan, J., & Barikdar, C. R. (2026). Embedding cybersecurity risk management into information systems governance: A national-scale framework for organizational resilience, economic stability, and critical infrastructure protection. International Journal of Applied Mathematics, 39(1s). https://doi.org/10.12732/ijam.v39i1s.1623
13. Hasan, S. N., Kaur, H., Mohonta, S. C., Siddiqa, K. B., Kaur, J., Haldar, U., et al. (2025a). The influence of artificial intelligence on data system security. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3476
14. Hasan, S. N., et al. (2025b). Self-healing cybersecurity systems using RL agents. In Proceedings of the 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11452866
15. Hassan, J., et al. (2025). Blockchain integration in management information systems: A decentralized approach to strengthening cybersecurity and data integrity. In Proceedings of the 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–7). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472020
16. Kaur, J., Prabha, M., Samiun, M., Hasan, S. N., Hasan, R., & Esa, H. (2025). Comparative analysis of transformer and LSTM architectures for cybersecurity threat detection using machine learning. EAI Endorsed Transactions on AI and Robotics, 4. https://publications.eai.eu/index.php/airo/article/view/9759
17. Khair, F. B., Saimon, A. S. M., Hossain, S., et al. (2025). Deep neural network-based imaging system for efficient pancreatic tumor identification. Intelligent Decision Technologies, 19(6), 4118–4129. https://doi.org/10.1177/18724981251381581
18. Mahin, M. R. H., Chakraborty, P., Das, N., Kaur, H., Himel, H. U., Kaur, J., & Mohapatra, A. G. (2026). Secured and standardized intelligent zero-touch 6G framework for edge-AI applications. IEEE Communications Standards Magazine. https://doi.org/10.1109/MCOMSTD.2026.3660159
19. Manik, M. M. T. G., Saimon, A. S. M., Ahmed, M. K., Hossain, S., Moniruzzaman, M., & Islam, M. S. (2025). Predictive modelling for early detection of type 2 diabetes using AI-driven machine learning algorithms and big data analytics. In J. C. Bansal, P. Jamwal, & S. Hussain (Eds.), Proceedings of the International Conference on AI and Robotics (AIR 2025) (Lecture Notes in Networks and Systems, Vol. 1629). Springer. https://doi.org/10.1007/978-3-032-05548-4_33
20. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 1273–1282). PMLR.
21. Mohonta, S. C., et al. (2026). Efficient homomorphic encryption techniques for confidential big-data analytics over public cloud infrastructures: Algorithms, performance, and trade-offs. In G. N. Nguyen, A. Swaroop, & P. Shukla (Eds.), Proceedings of Fifth International

Conference on Computing and Communication Networks (ICCCN 2025) (Lecture Notes in Networks and Systems, Vol. 1859). Springer. https://doi.org/10.1007/978-3-032-21499-7_41

22. Morgan, S. (2020). Cybercrime to cost the world $10.5 trillion annually by 2025. Cybercrime Magazine. https://cybersecurityventures.com/hackerpocalypse-cybercrime-report-2016/
23. Nabi, N., Tuhin, M. K., Bashir, M., Lucky, K. Y., Raihan, M., & Imam, H. (2026). Continual learning pipelines for detecting insider threats across corporate cybersecurity infrastructures. In Proceedings of the 2025 International Conference on Electrical Engineering and Informatics (ICEEI) (pp. 1–8). IEEE. https://doi.org/10.1109/ICEEI68459.2025.11330487
24. Nusrat, S., Murshed, H., & Afrin, S. (2025). Antibiotic resistance at the human–animal–environment crossroads: A systematic review of the silent global pandemic. Microbial Bioactives, 8(1), 1–7. https://doi.org/10.25163/microbbioacts.8110426
25. Orthi, S. M., Chakraborty, P., Siam, M. A., Shan-A-Alahi, A., Al Zaiem, A., Hasan, S. N., Kaur, J., Mahmud, F., & Goffer, M. A. (2023). AI-driven cyber threat intelligence as a management information system: Integrating cybersecurity governance and IT project management for organizational resilience. The Eastasouth Journal of Information System and Computer Science, 1(2), 194–214. https://doi.org/10.58812/esiscs.v1i02.873
26. Orthi, S. M., Rahman, M. H., Siddiqa, K. B., Uddin, M., Hossain, S., Mamun, A. A., & Khan, M. N. (2025a). Federated learning with privacy-preserving big data analytics for distributed healthcare systems. Journal of Computer Science and Technology Studies, 7(8), 269–281. https://doi.org/10.32996/jcsts.2025.7.8.31
27. Orthi, S. M., et al. (2025b). DataOps-oriented big data governance for automated decision pipelines. In Proceedings of the 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11452860
28. Raihan, M., Adnan, M., Hossain, M. J., Siddiqa, K. B., Karim, F., & Mohonta, S. C. (2026). Adversarial robustness mechanism for safeguarding biometric verification across mobile financial applications. In Proceedings of the 2025 International Conference on Electrical Engineering and Informatics (ICEEI) (pp. 1–7). IEEE. https://doi.org/10.1109/ICEEI68459.2025.11330502
29. Rozario, E., et al. (2025). Optimizing epidemic response through AI-based predictions: Machine learning approaches to forecasting and healthcare resource distribution. In Proceedings of the 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–6). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472546
30. Shan-A-Alahi, A., et al. (2026). Deep learning-based threat prediction and autonomous response mechanisms for containerized microservices in hybrid cloud deployments. In G. N. Nguyen, A. Swaroop, & P. Shukla (Eds.), Proceedings of Fifth International Conference on Computing and Communication Networks (ICCCN 2025) (Lecture Notes in Networks and Systems, Vol. 1859). Springer. https://doi.org/10.1007/978-3-032-21499-7_42
31. Siam, M. A., Lucky, K. Y., Hasan, S. N., Kaur, J., Kaur, H., Uddin, M. S., & Manik, M. M. T. G. (2025a). Cybersecure intelligent sensor framework for smart buildings: AI-based intrusion detection and resilience against IoT attacks. Sensors, 25(24), 7680. https://doi.org/10.3390/s25247680
32. Siam, M. A., Shan-A-Alahi, A., Tuhin, M. K., Hossain, E., Bashir, M., Lucky, K. Y., & Al Zaiem, A. (2025b). AI-driven cyber threat intelligence systems: A national framework for proactive defense against evolving digital warfare. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3793
33. Siddiqa, K. B., Rahman, H., Barikdar, C. R., Orthi, S. M., Miah, M. A., & Manik, M. M. T. G. (2025). Assessment of survivability and importance analysis for networks managing intricate traffic flows. IEEE Communications Standards Magazine. https://doi.org/10.1109/MCOMSTD.2025.3638981
34. Siddiqa, K. B., Rahman, H., Barikdar, C. R., Orthi, S. M., Miah, M. A., & Rahman, R. (2024). AI-driven project management systems: Enhancing IT project efficiency through MIS integration. In Proceedings of the 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 114–119). IEEE. https://doi.org/10.1109/ICPIDS65698.2024.00027
35. Uddin, S. M. M., et al. (2025). Bio-cognitive AI systems for predictive healthcare decision support. In Proceedings of the 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–9). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11453175
36. Varanasi, S. R., Valiveti, S. S. S., Adnan, M., Faruk, M. I., Hossain, M. J., & Manik, M. M. T. G. (2026). Cross-domain standardization and secure edge intelligence for real-time digital twin deployments in next-generation communication systems. IEEE Communications Standards Magazine. https://doi.org/10.1109/MCOMSTD.2026.3662187
37. Zerine, I., Islam, M. M., Khan, M. A. U., Chy, M. A. R., Saimon, A. S. M., Manik, M. M. T. G., & Wata, C. (2026). Explainable churn prediction in telecom with tabular ML five model benchmark and SHAP analysis. Discover Artificial Intelligence. https://doi.org/10.1007/s44163-026-00983-0